IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Appellants: Robert Jochemsen et al. Group Art Unit: 2186

Serial No.: 10/533,735 Examiner: Schnee, Hal W.

Filed: May 4, 2005 Confirmation No.: 4556

For: METHOD AND DEVICE FOR PERSISTENT-MEMORY MANAGEMENT

Mail Stop Appeal Brief-Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

REPLY BRIEF UNDER 37 C.F.R. § 41.41(a)

This is an appeal to the Board of Patent Appeals and Interferences from the decision of the Examiner dated August 7, 2008, which finally rejected claims 1-8 and 10-25 in the above-identified application. An Appeal Brief was filed on January 7, 2009. This Reply Brief is in response to the Examiner's Answer dated March 25, 2009. This Reply Brief is hereby submitted pursuant to 37 C.F.R. § 41.41(a).

CERTIFICATE OF MAILING UNDER 37 C.F.R. 1.8

I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being facsimile transmitted to the Patent and Trademark Office on the date shown below.

Date:_____ Signed:_____

Typed Name: Mark A. Wilson

TABLE OF CONTENTS

I. STATUS OF CLAIMS

Claims 1-8 and 10-25 are pending in the present application.

No claims are objected to.

Claim 9 is canceled.

No claims are withdrawn.

No claims are allowed.

Claims 1-8 and 10-25 stand rejected, as follows:

Claims 1-4, 7, 8, and 10 were rejected under 35 U.S.C. 103(a) as being unpatentable over Stockdale et al. (U.S. Pat. No. 6,804,763, hereinafter Stockdale) in view of Cheng et al. (U.S. Pat. No. 5,701,516, hereinafter Cheng) and further in view of O'Neill (U.S. Pat. Pub. No. 2003/0182414, hereinafter O'Neill).

Claims 5, 6, and 11-15 were rejected under 35 U.S.C. 103(a) as being unpatentable over Stockdale in view of Cheng further in view of O'Neill and further in view of Hanes (U.S. Pat. Pub. No. 2003/0081932, hereinafter Hanes).

Claims 16-25 were rejected under 35 U.S.C. 103(a) as being unpatentable over Stockdale in view of Cheng in view of O'Neill in view of Hanes and further in view of Lee et al. (U.S. Pat. No. 5,903,167, hereinafter Lee).

Claims 1-8 and 10-25 are the subject of this appeal. A copy of claims 1-8 and 10-25 is set forth in the Claims Appendix.

II. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Whether claims 1-4, 7, 8, and 10 are patentable over the combination of Stockdale, Cheng, and O'Neill under 35 U.S.C. 103(a).

B. Whether claims 5, 6, and 11-15 are patentable over the combination of Stockdale, Cheng, O'Neill, and Hanes under 35 U.S.C. 103(a).

C. Whether claims 16-25 are patentable over the combination of Stockdale, Chang, O'Neill, Hanes, and Lee under 35 U.S.C. 103(a).

III. ARGUMENT

The following remarks should be considered within the context of the arguments presented in Appellants' Appeal Brief and the Examiner's Response. For clarity, the arguments presented herein are arranged in sections and subsections which correspond to the sections and subsections of the Appeal Brief. Where additional explanation is warranted, in light of the Examiner's Response, to elaborate on the Arguments present in the Appeal Brief, those additional remarks are presented below.

A. Claims 1-4, 7, 8, and 10 are patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of references does not teach all the limitations of the claims.

Appellants respectfully submit that claim 1 is patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of references does not teach all of the limitations of the claim. Claim 1 is recites:

> A memory management device for managing a memory space of at least one persistent-memory device, comprising:
>
> a memory allocation unit adapted to communicate with at least one application device and to allocate at least one first part of said memory space to said application device to write a first working data structure comprising a plurality of working data blocks to the memory space and to write a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure, wherein said allocation unit is further adapted to communicate with at least one file system device, and to allocate on request from said application device or from said file system device said first part of said memory space to said file system; wherein the memory space is used as a write cache memory for said file system device.

(Emphasis added.)

As stated in the Appeal Brief, the combination of Stockdale, Cheng, and O'Neill does not teach all of the limitations of the claim. In particular, the combination of cited references does not teach writing a second working data structure which includes a copy of the plurality of working data blocks, in which the second working data structure includes a copy of the first working data structure in the same memory space as the first

working data structure, as recited in the claim. Additionally, the combination of cited references does not teach an allocation unit which is adapted to communicate with at least one file system device, and to allocate on request from the application device or from the file system device the first part of the memory space to the file system. The remarks within this Reply Brief focus on the first argument that the combination of cited references does not teach first and second working data structures in the same memory space.

1. O'Neill does not teach writing a second working data structure in the same memory space as the first working data structure.

The Examiner relies on O'Neill as purportedly teaching writing a second working data structure in the same memory space as the first working data structure. In particular, the Examiner relies on the description of the backup bank 1234 and the original first bank 1242 in the flash memory. However, even though the backup bank 1234 and the original first bank 1242 are stored in the same flash memory, O'Neill does not teach writing a second working data structure in the same memory space as the first working data structure at least because the backup bank 1234 is not a working data structure.

O'Neill refers to the backup bank 1234 as a copy of the working bank 1232, but the designation as a copy of the working bank nevertheless fails to make the backup bank a working bank. The working bank 1232 is specifically allocated in the RAM memory, rather than the flash memory, and is used within the working environment for the update agent 1025. O'Neill, paragraph 145, lines 1-4. More specifically, the working bank 1232 acts as an operational buffer or working memory area where operations determined by the instruction set are performed. O'Neill, paragraph 145, lines 5-9. In contrast, the backup bank merely serves as a backup copy of the working bank and is only used for fault tolerance in case of power interruptions or corruption of the working bank. Thus, the backup bank is not a working bank because, unlike the working bank, the backup bank is not used for operations or instructions.

Therefore, even if the original first bank 1242 and the backup bank 1234 were considered to be stored within the same memory space because they are both stored in the flash memory, the original first bank 1242 and the backup bank 1234 are not first and

second working data structures because the backup bank is not a working bank, accordingly to the explicit designations used in O'Neill. In other words, O'Neill does not teach first and second working data structures stored in the same memory space at least because the backup bank of O'Neill is not a working data structure. Since the backup bank of O'Neill is not a working data structure, within the context of the claims of the present application, the Examiner's reliance on the teachings of O'Neill is misplaced because O'Neill does not teach first and second working data structures stored within the same memory space. Accordingly, the combination of cited references relied on by the Examiner do not teach all of the limitations of the claim because O'Neill does not teach first and second working data structures stored within the same memory space.

Therefore, the combination of Stockdale, Cheng, and O'Neill does not teach the limitations of claim 1 because the combination of Stockdale, Cheng, and O'Neill does not teach writing a second working data structure in the same memory space as the first working data structure. Accordingly, Appellants respectfully assert claim 1 is patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of Stockdale, Cheng, and O'Neill does not teach all of the limitations of the claim. Consequently, Appellants request that the rejection of claim 1 under 35 U.S.C. 103(a) be withdrawn.

Appellants respectfully assert independent claim 7 is also patentable over Stockdale, Cheng, and O'Neill at least for similar reasons to those stated above in regard to the rejection of independent claim 1. In particular, the rejection of claim 7 relies on similar reasoning to characterize the writing of a second working data structure in the same memory space as the first working data structure. Here, although the language of claim 7 differs from the language of claim 1, and the scope of claim 7 should be interpreted independently of claim 1, Appellants respectfully assert that the remarks provided above in regard to claim 1 also apply to the rejection of claim 7. Accordingly, Appellants respectfully assert independent claim 7 is also patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of Stockdale, Cheng, and O'Neill does not teach all the limitations of the claim. Consequently, Appellants request that the rejection of claim 7 under 35 U.S.C. 103(a) be withdrawn.

Appellants respectfully assert independent claim 10 is also patentable over Stockdale, Cheng, and O'Neill at least for similar reasons to those stated above in regard to the rejection of independent claim 1. In particular, the rejection of claim 10 relies on similar reasoning to characterize the writing of a second working data structure in the same memory space as the first working data structure. Here, although the language of claim 10 differs from the language of claim 1, and the scope of claim 10 should be interpreted independently of claim 1, Appellants respectfully assert that the remarks provided above in regard to claim 1 also apply to the rejection of claim 10. Accordingly, Appellants respectfully assert independent claim 10 is also patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of Stockdale, Cheng, and O'Neill does not teach all the limitations of the claim. Consequently, Appellants request that the rejection of claim 10 under 35 U.S.C. 103(a) be withdrawn.

Given that claims 2-4 and 8 depend from and incorporate all of the limitations of the independent claims 1 and 7, which are patentable over the cited references, Appellants respectfully submit that dependent claims 2-4 and 8 are also patentable over the cited reference based on allowable base claims. Additionally, each of claims 2-4 and 8 may be allowable for further reasons. Consequently, Appellants request that the rejections of claims 2-4 and 8 under 35 U.S.C. 103(a) be withdrawn.

2. Stockdale does not teach an allocation unit that is adapted to communicate with at least one file system.

The remarks of the Examiner's Response do not necessitate further clarification. Appellants respectfully submit that the arguments presented in the Appeal Brief are sufficient to distinguish the indicated language of the claims from the teachings of Stockdale within the proposed combination of Stockdale, Cheng, and O'Neill.

B. Claims 5, 6, and 11-15 are patentable over Stockdale, Cheng, O'Neill, and Hanes because the combination of references does not teach all the limitations of the claims.

The remarks of the Examiner's Response do not necessitate further clarification. Appellants respectfully submit that the arguments presented in the Appeal Brief are

sufficient to distinguish the indicated language of the claims from the teachings of the proposed combination of Stockdale, Cheng, O'Neill, and Hanes.

C. Claims 16-25 are patentable over the combination of Stockdale, Cheng, O'Neill, Hanes, and Lee because the combination of references does not teach all of the limitations of the claims.

The remarks of the Examiner's Response do not necessitate further clarification. Appellants respectfully submit that the arguments presented in the Appeal Brief are sufficient to distinguish the indicated language of the claims from the teachings of the proposed combination of Stockdale, Cheng, O'Neill, Hanes, and Lee.

IV. CONCLUSION

For the reasons stated above, claims 1-8 and 10-25 are patentable over the cited references. Thus, the rejections of claims 1-8 and 10-25 should be withdrawn. Appellants respectfully request that the Board reverse the rejections of claims 1-8 and 10-25 under 35 U.S.C. 103(a) and, since there are no remaining grounds of rejection to be overcome, direct the Examiner to enter a Notice of Allowance for claims 1-8 and 10-25.

At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account **50-4019** pursuant to 37 C.F.R. 1.25. Additionally, please charge any fees to Deposit Account **50-4019** under 37 C.F.R. 1.16, 1.17, 1.19, 1.20 and 1.21.

Respectfully submitted,

/mark a. wilson/

Date: May 26, 2009

Mark A. Wilson
Reg. No. 43,994

Wilson & Ham
PMB: 348
2530 Berryessa Road
San Jose, CA 95132
Phone: (925) 249-1300
Fax: (925) 249-0111

V. CLAIMS APPENDIX

1. A memory management device for managing a memory space of at least one persistent-memory device, comprising:

a memory allocation unit adapted to communicate with at least one application device and to allocate at least one first part of said memory space to said application device to write a first working data structure comprising a plurality of working data blocks to the memory space and to write a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure, wherein said allocation unit is further adapted to communicate with at least one file system device, and to allocate on request from said application device or from said file system device said first part of said memory space to said file system; wherein the persistent memory is used as a write cache memory for said file system device.

2. A memory management device according to claim 1, wherein said memory allocation unit is adapted to maintain a memory allocation table at a current status, said memory allocation table assigning at least one memory address representing a defined part of said memory space to either said application device or to said file system device.

3. A memory management device according to claim 2, further comprising a processor and a memory, wherein said memory allocation unit is implemented in the form of at least one first executable file contained in said memory.

4. A memory management device according to claim 3, wherein said memory is a persistent-memory device, in particular said persistent-memory device.

5. A file system device comprising a file allocation unit adapted to maintain a file allocation table at a current status, said file allocation table assigning at least one disk space address to at least one file, wherein said file allocation unit is adapted to communicate with a memory management device that is related to a persistent-memory

device and to include an address of at least one first memory space of said persistent-memory device in the maintenance of said file allocation table, wherein the file system device is configured to convert a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium, wherein the copy of the first working data structure is written to a same memory space as the first working data structure, and wherein the persistent memory is used as a write cache memory for said file system device.

6. A file system device according to claim 5, further comprising a processor and a memory, wherein said file allocation unit is implemented in the form of at least one second executable file contained in said memory.

7. An application device, comprising a persistent-memory device connected to a processor, and a data management unit adapted to manipulate data in said persistent memory device, wherein said data management unit is adapted to write at least one third executable file to said persistent memory device, or to provide the file system with a reference to at least one third executable file in said file system, such that by executing said third executable file said processor is adapted to transform a copy of a first working data structure into a predetermined data-sequence form, wherein the copy of the first working data structure is stored in a same memory space as the first working data structure; and wherein the persistent memory is used as a write cache memory for a file system device.

8. An application device according to claim 7, wherein said data management unit is provided in the form of least one fourth executable file in a memory, particularly, in said persistent memory.

9. (canceled)

10. A data processing system, comprising a memory management device for managing a memory space of at least one persistent-memory device, comprising a

memory allocation unit adapted to communicate with at least one application device and to allocate at least one first part of said memory space to said application device to write a first working data structure comprising a plurality of working data blocks to the memory space and to write a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure, wherein said allocation unit is further adapted to communicate with at least one file system device, and to allocate on request from said application device or from said file system device said first part of said memory space to said file system; wherein the persistent memory is used as a write cache memory for said file system device.

11. A method for managing memory space of a persistent-memory device, comprising:

allocating at least one first part of said memory space to a file system device;

writing a first working data structure comprising a plurality of working data blocks to the memory space;

writing a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure; and

converting the second working data structure into a predetermined data-sequence form;

wherein the persistent memory is used as a write cache memory for said file system device.

12. A method according to claim 11, wherein said allocating step comprises a step of blocking a writing access to said first part of said memory space.

13. A method according to claim 12, wherein said allocating step comprises a step of giving away to said file system device the power of reading access to said first part of said memory space.

14. A method according to claim 11, comprising a step of deallocating said first part of said memory space to a memory management device.

15. A method according to claim 14, wherein said allocating step or said deallocating step comprises transmitting an address range defining said first part of said memory space from said memory management device to said file system device or, respectively, vice versa.

16. A method according to claim 14, wherein said deallocating step is performed for said first part of said memory space given the condition that first data contained in said first part of said memory space is stored in the form of file data in a second part of said memory space, said file data having a predetermined file structure, and that said second part of said memory space is allocated to said file system device.

17. A method according to claim 16, wherein said deallocating step is performed for said second part of said memory space given the condition that said file data has been written to a secondary storage medium.

18. A method for write-caching first data worked on by an application, said first data being contained in a first part of a memory space of a persistent-memory device, comprising a step of performing a memory managing method according to claim 17.

19. A write-caching method according to claim 18, comprising, after said allocating step, a step of sending a confirmation message from said file system device to said application device.

20. A write-caching method according to claim 18, wherein said first data is a copy of third data contained in a third part of said memory space, said write-caching method comprising, before performing said memory managing method, a step of copying said third data to said first memory space.

21. A write-caching method according to claim 18, comprising the steps of

allocating a fourth part of said memory space to said application device for an executable file or dynamic link library that is adapted to converting said first data into file data

writing said executable file or dynamic link library to said fourth part of said memory space

allocating said fourth part of said memory space to said file system device.

22. A write-caching method according to claim 21, comprising a step of transforming said first data into said file data with the aid of said executable file or said dynamic link library.

23. A write-caching method according to claim 22, wherein said transforming step is initiated by said file system device.

24. A write-caching method according to claim 23, comprising, after said transforming step, a step of deallocating said fourth part of said memory space to said memory management device.

25. A method for saving data worked on by an application device to a file on a secondary storage medium, comprising performing a write-caching method according to claim 18, and further comprising writing said file data to said secondary storage medium.